



Quantum Time-Space Tradeoffs for Sorting *

Hartmut Klauck[†]
 School of Mathematics
 Institute for Advanced Study
 Princeton, NJ08540, USA
 klauck@ias.edu

ABSTRACT

We investigate the complexity of sorting in the model of sequential quantum circuits. While it is known that in general a quantum algorithm based on comparisons alone cannot outperform classical sorting algorithms by more than a constant factor in time complexity, this is wrong in a space bounded setting. We observe that for all storage bounds $n/\log n \geq S \geq \log^3 n$, one can devise a quantum algorithm that sorts n numbers (using comparisons only) in time $T = O(n^{3/2} \log^{3/2} n / \sqrt{S})$. We then show the following lower bound on the time-space tradeoff for sorting n numbers from a polynomial size range in a general sorting algorithm (not necessarily based on comparisons): $TS = \Omega(n^{3/2})$. Hence for small values of S the upper bound is almost tight. Classically the time-space tradeoff for sorting is $TS = \Theta(n^2)$.

Categories and Subject Descriptors

F.2 [Theory of Computation]: analysis of algorithms and problem complexity

General Terms

Theory, Algorithms

Keywords

quantum algorithms, sorting, lower bounds, complexity theory

1. INTRODUCTION

Sorting is arguably one of the most important and well-studied problems in computer science. While any comparison-based algorithm needs $\Omega(n \log n)$ time to sort n numbers, and several algorithms achieving a matching running

time are well known, the situation changes if we are confronted with the problem to sort when the dataset is too large to fit into the memory of our computer. In this setting we are given a bound $S(n)$ on the available memory and have to sort with an algorithm whose space requirements do not exceed this bound. This situation arises e.g. if datasets are stored distributed on the internet and are too large to be held in the memory of the computer at a single site.

After prior results by Borodin and Cook in [7], Beame [3] proved a lower bound of $TS = \Omega(n^2)$ for any classical algorithm sorting n numbers. A matching upper bound for sorting on a RAM has been proved in [16], for all space bounds $\log n \leq S \leq n/\log n$. This upper bound is actually achieved by an algorithm that accesses the data only using comparisons. This is remarkable, since sometimes operations on the numbers itself can speedup sorting, e.g. Radixsort (see e.g. [11]) can sort polynomial size numbers in linear time, while comparison based algorithms need time $\Omega(n \log n)$ for this task. So the classical time-space complexity of sorting is well known. Interestingly, while Beame [3] investigates the UNIQUE ELEMENTS problem (output the list of elements appearing exactly once in the input), Borodin and Cook [7] consider the RANKING problem (output the permutation needed to sort) and show $TS = \Omega(n^2/\log n)$ for numbers from a quadratic size range. No tight analysis of RANKING seems to have appeared since (only a modest improvement in [17]), and the UNIQUE ELEMENTS problem is not fit to provide lower bounds in the situation when all the inputs are known to be distinct, which they are with high probability when drawn uniformly from a large range.

Quantum computing is an active research area offering interesting possibilities to obtain improved solutions to information processing tasks by employing computing devices based on quantum physics, see e.g. [15] for a nice introduction into the field. It has been shown by Høyer et al. in [14], however, that any comparison-based quantum sorting algorithm needs time $\Omega(n \log n)$, which is already achieved by classical algorithms. So is there no point in using quantum computers to sort? We demonstrate that in the space bounded setting the quantum complexity of sorting is quite different from the classical complexity.

We use the model of quantum circuits to model time-space tradeoffs. While in the classical setting branching programs are the standard model to investigate these problems, employing quantum branching programs seems to complicate the definition of our model unnecessarily. A quantum circuit uses space S , if it operates on S working qubits. Furthermore it has write-only access to several output qubits, since

*A full version of this paper is available as quant-ph/0211174.

[†]Supported by NSF grant CCR-9987845.

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. To copy otherwise, to republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee.

STOC'03, June 9–11, 2003, San Diego, California, USA.
 Copyright 2003 ACM 1-58113-674-9/03/0006 ...\$5.00.

this output is typically too large to be held in the working memory during the whole computation. The circuit accesses the input by specific oracle gates. For the upper bound we use a comparison oracle, i.e., an oracle gate is fed (superpositions over) two indices of numbers to be compared and outputs a bit indicating the comparison result into an extra qubit. For the lower bound we choose to restrict the size of the numbers in the input to n^2 , and here we consider an access oracle, that directly reads numbers from the input into the work space.

Using Grover's famous quantum search algorithm and its adaptation to minimum finding by Dürr and Høyer in [12] we can show the following:

THEOREM 1. *For all S in $[\Omega(\log n), \dots, O(n/\log n)]$ there is a quantum circuit with space S that, given a comparison oracle for n numbers, outputs the sorted sequence, and uses time $O(n^{3/2} \log^{3/2} n / \sqrt{S})$. The entire output is correct with probability $1 - \epsilon$ for an arbitrarily small constant $\epsilon > 0$.*

If we restrict the size of the numbers, we can use the same algorithm while employing an access oracle.

COROLLARY 1. *For all S in $[\Omega(\log n), \dots, O(n/\log n)]$ there is a quantum circuit with space S that, given an access oracle for n numbers from $[1, \dots, n^k]$ for some constant k , outputs the sorted sequence in time $O(n^{3/2} \log^{3/2} n / \sqrt{S})$. The entire output is correct with probability $1 - \epsilon$ for an arbitrarily small constant $\epsilon > 0$.*

The quantum sorting algorithm clearly beats the classical time-space tradeoff lower bound. Actually it gives rise to a tradeoff $T^2 S = O(n^3 \log^3 n)$. It performs best whenever S is small. E.g. for $S = \log^k(n)$ the running time is $O(n^{3/2} / \log^{(k-3)/2} n)$, a clear improvement compared to the classical bound $\Omega(n^2 / \log^k n)$.

We then investigate the question whether one can do even better. First let us fix an output convention for sorting. Denote $\text{Min}(x, i)$ the i th number of the sorted sequence. Let us assume an algorithm outputs a sequence $(y_1, j_1), \dots, (y_n, j_n)$, so that for all i the element y_i is supposed to equal $\text{Min}(x, j_i)$, and that $\{j_1, \dots, j_n\} = \{1, \dots, n\}$. Since we consider quantum algorithms we will allow errors. Furthermore we will require that the j_i are produced in the same order for all inputs. Our main result is the following lower bound.

THEOREM 2. *Let A be any quantum circuit that, given an access oracle for a sequence x of n numbers from $[1, \dots, n^2]$, outputs n pairs of numbers $O_1(x), \dots, O_n(x)$, so that the probability that $O_i(x) = (\text{Min}(x, j_i), j_i)$ is $2/3$ for each i and the j_i are a fixed permutation of $\{1, \dots, n\}$. Suppose A uses S work qubits and T oracle gates, then $ST = \Omega(n^{3/2})$.*

Note that we do not even require that the whole sorted sequence is ever simultaneously output correctly.

Hence for small S we cannot substantially speed up the quantum sorting algorithm. E.g. for $S = \text{poly}(\log n)$ the necessary and sufficient time to sort on a quantum computer is $\tilde{\Theta}(n^{3/2})$.

Technically the lower bound proof proceeds as follows. A given quantum circuit is sliced into parts containing only $\delta\sqrt{n}$ queries. Each such slice starts with some initial information stored in the S qubits that has been computed by the previous slices. We show how to get rid of this initial

information while deteriorating the success probability only exponentially in S . This step is similar to an application of the union bound (classically the success probability for one of the 2^S fixed states of the S bits given in the beginning must be $1/2^S$ times as good as the success probability with initial information). We show that something similar is possible in the quantum case.

Then we are left with the problem to analyze the success probability of a slice of the circuit without initial information, and show that it can be only large enough, if the number of outputs is small. Since there have to be n outputs this leads to a lower bound on the number of slices and hence the time complexity.

2. DEFINITIONS AND PRELIMINARIES

In this section we give some background on quantum states and their distinguishability, define the model of quantum circuits we are studying, describe a quantum version of the union bound, and discuss lower bounds on query complexity by query magnitude arguments. For more quantum background we refer to [15].

2.1 Quantum States

The quantum mechanical analogue of a random variable is a probability distribution over superpositions, also called a *mixed state*. For the mixed state $X = \{p_i, |\phi_i\rangle\}$, where $|\phi_i\rangle$ has probability p_i , the *density matrix* is defined as $\rho_X = \sum_i p_i |\phi_i\rangle\langle\phi_i|$. Density matrices are Hermitian, positive semi-definite, and have trace 1. I.e., a density matrix has an eigenvector basis, all the eigenvalues are real and between zero and one, and they sum up to one.

The *trace norm* of a matrix A is defined as $\|A\|_t = \text{Tr} \sqrt{A^\dagger A}$, which is the sum of the magnitudes of the singular values of A . Note that if ρ is a density matrix, then it has trace norm one.

A useful theorem states that for two mixed states ρ_1, ρ_2 their distinguishability is reflected in $\|\rho_1 - \rho_2\|_t [1]$:

FACT 1. *Let ρ_1, ρ_2 be two density matrices on the same space $\mathcal{H}$. Then for any measurement $\mathcal{O}$,*

$$\left\| \rho_1^{\mathcal{O}} - \rho_2^{\mathcal{O}} \right\|_1 \leq \|\rho_1 - \rho_2\|_t,$$

where $\rho^{\mathcal{O}}$ denotes the classical distribution on outcomes resulting from the measurement of ρ , and $\|\cdot\|_1$ is the ℓ_1 norm. Furthermore, there is a measurement $\mathcal{O}$, for which the above is an equality.

2.2 Quantum Circuits

We now define quantum oracle circuits. The two parameters we are interested in are the number of work qubits corresponding to the space bound, and the number of queries, which is always smaller than the overall number of gates, corresponding to the time.

A *quantum circuit* on S work qubits and M output qubits is defined as an ordered set of gates, where each gate consists of a unitary operation on some k qubits and a specification of these k qubits. The gates either operate on the work qubits only (in which case we can assume that $k = S$), or they are output gates, that perform a controlled-not, where the control is among the work qubits, and the target qubit is an output qubit. We require that each output qubit is used

only once. In this way the output qubits are usable only to record the computation results.

Another type of gates are *query gates*, and they are needed to access the input. In general a query gate performs the following operation for some input $x \in \{1, \dots, 2^k\}^n$:

$$|i\rangle|a\rangle \mapsto |i\rangle|a \oplus x_i\rangle,$$

for all classical strings i, a , where the length of i is $\log n$ and the length of a is k . The behavior of a query gate on superpositions is defined by linearity. In a *computation* all qubits start blank, and then the gates are applied in the order defined on them. For sorting we consider two different types of input oracles. In a *comparison* oracle the algorithm is allowed to query an n^2 -bit string that contains the results of the pairwise comparisons between n numbers. In an *access* oracle to input $x_1, \dots, x_n$ the query gate directly reads numbers x_i (or their superpositions). Note that an access oracle can efficiently simulate a comparison oracle if the numbers to be sorted are from a polynomial size range.

We require that a sorting circuit makes outputs in the following way: if the input is presented as a comparison oracle, then the algorithm is supposed to output the inverse of the permutation needed to sort. If the input is given as an access oracle, then the output is given as a list of numbers with their positions. We require that a number is output as a whole at some point in the algorithm, i.e., there are k one bit output gates directly following each other when a number from a 2^k range is output.

Let us consider the output convention more closely (in the case of an access oracle). For some input x let O_i denote the i th output, and $O_i(x)$ the probability distribution obtained by measuring this output. This is a distribution on pairs (y_i, j_i) . Let $\text{Min}(x, i)$ denote the position of the i th element of the sorted sequence in the input, i.e., the pair $(\text{Min}(x, j_i), j_i)$ is supposed to be produced at output O_i . The sequence j_i has to be some fixed permutation of $\{1, \dots, n\}$ with probability 1.

We say that $O_i(x)$ is correct with probability $1 - \epsilon$, if $O_i(x) = (\text{Min}(x, j_i), j_i)$ with probability $1 - \epsilon$. We require in our lower bound theorem that each $O_i(x)$ is correct with probability $2/3$.

A circuit is said to have *time complexity* T , if the number of oracle gates is T .

2.3 A Quantum Union Bound

In this subsection we want to develop a simple tool needed to take away the initial information in a slice of the sorting circuit.

Let ρ denote the density matrix of a state on m qubits. We want to replace ρ by the completely mixed state, while retaining some of the success probability of an algorithm taking ρ as an input. The following lemma will be helpful.

LEMMA 2. *Let ρ be any density matrix on m qubits. Let M denote the density matrix of the completely mixed state, i.e., the matrix with entries $1/2^m$ on the diagonal and zeros elsewhere. Then there exists a density matrix σ , so that*

$$M = 1/2^m \rho + (1 - 1/2^m) \sigma.$$

PROOF. We have to show that $\sigma = (M - \rho/2^m)/(1 - 1/2^m)$ is a density matrix. For ρ being a density matrix, σ is clearly Hermitian and has trace 1. So we have to show that σ is positive semidefinite, which is equivalent to $I - \rho$ being

positive semidefinite, for the identity matrix I . Let U be some unitary transformation that diagonalizes ρ , i.e., $D = U\rho U^\dagger$ is diagonal. U exists, since ρ is Hermitian. Clearly $I - \rho$ is positive semidefinite iff $I - D$ is, and D contains on its diagonal nonnegative numbers that sum to 1, hence $I - D \geq 0$. $\square$

We want to apply the above lemma in the following way.

LEMMA 3. *Suppose there is an algorithm that on some input x first receives S qubits of initial information depending arbitrarily on x , and that makes afterwards only queries to the input. Suppose the algorithm produces some output correctly with probability p .*

Then there is an algorithm that uses no initial information, makes the same number of queries, and has success probability $p/2^S$.

Actually the above lemma can be thought of as a quantum union bound. Note that the S qubits can be in as many states as there are inputs x , still removing them decreases the success probability by a factor exponentially in S only.

PROOF. Replace the quantum state containing the initial information by the completely mixed state M on S qubits. Then run the algorithm in exactly the same manner as before. Clearly the algorithm does not get initial information in this way. Due to Lemma 2 the original state has some probability $1/2^S$ in M , respectively one can view M as a mixture of the original state with probability $1/2^S$ and another state with the remaining probability. So also the outcome of the algorithm is such a mixture, and if the success probability was p originally, it must be at least $p/2^S$ for the modified algorithm. $\square$

Note that the completely mixed state on S qubits can easily be obtained from a blank state on $2S$ qubits by performing Hadamard gates on the qubits 1 to S and then controlled not gates on the pairs $i, S + i$.

Another way of obtaining a similar result is to use quantum teleportation (invented by [6], see also [15]). In the teleportation scheme two players who share m EPR-pairs can communicate an arbitrary quantum state in the following way. If player Alice holds a quantum state ρ on m qubits, she applies measurements in the Bell basis to the m pairs of qubits given by one qubit from an EPR-pair and one qubit of ρ each. Bob, holding the m other qubits belonging to the EPR-pairs then gets a message from Alice containing her measurement results in $2m$ classical bits. He is then able to perform certain operations on his qubits depending on the message, which enable him to recover ρ .

It is known that for each ρ , the probability of each of the possible measurement results is exactly 4^{-m} . Furthermore for one of the measurement results, Bob does not have to do anything to his qubits to get ρ , i.e., with probability 4^{-m} the state of Bob is correct already. Note that this implies that Bob's qubits before he receives the message, which are in a completely mixed state, can be viewed as an ensemble of states $\rho_1, \dots, \rho_{4^m}$, where $M = \sum_i 4^{-m} \rho_i$, and $\rho_1 = \rho$.

2.4 Lower Bounds for Query Algorithms

In [5] an $\Omega(\sqrt{n})$ lower bound for the problem of finding a marked element in an unordered database of size n has been given, matching the upper bound of Grover's algorithm [13].

This lower bound relies on the notion of *query magnitude*. For other lower bound techniques for query complexity see e.g. [10].

The query magnitude technique is basically an adversary argument. An adversary is able to change the black-box input without the query algorithm noticing that (for a more refined type of quantum adversary argument see [2]). We use the following statement derived as Corollary 3.4 in [5].

FACT 4. *Let $x = x_1, \dots, x_n$ be an input given as an access oracle, with x_i from some finite set, and let $x'(i)$ be any input that differs from x in position i and nowhere else. A is any quantum algorithm that accesses the oracle via at most T queries. The state ρ_x denotes the state of A 's workspace when querying x , the state $\rho_{x'(i)}$ when querying $x'(i)$.*

Then for any $\alpha > 0$ there is a set of at least $n - T^2/\alpha^2$ input positions i such that for all $x'(i)$: $\|\rho_x - \rho_{x'(i)}\|_t \leq 2\alpha$.

In our lower bound we will need a somewhat stronger statement that allows us to deal with a situation conditioned on results of previous measurements.

LEMMA 5. *Let $x = x_1, \dots, x_n$ be an input given as an access oracle, with x_i from some finite set, and let $x'(i)$ be any input that differs from x in position i and nowhere else. A is any quantum algorithm that accesses the oracle via at most T queries. Suppose A contains no measurements, but at the end a measurement is performed on some of the qubits. Fix some outcome F of this measurement that occurs with probability q_x . Then assume that some event E happens with probability p_x conditional on F .*

Then for any $\alpha > 0$ there is a set of at least $n - T^2/\alpha^2$ input positions i such that if A is performed on $x'(i)$, then the probability that F is the outcome of the measurement and E happens is at least $q_x(p_x - 2\alpha)$.

PROOF. Let U_x and $U_{x'(i)}$ denote the unitary transformations done by the circuit A on inputs x and $x'(i)$. W.l.o.g. A starts from the blank state $|0\rangle$ on some qubits. Let $|\phi_x\rangle$ denote the state obtained by performing $U_x|0\rangle$, and measuring the resulting state, assuming that F is the result of the measurement. Let $|\psi_x\rangle = U_x^{-1}|\phi_x\rangle$. Clearly there is a state $|\theta_x\rangle$ so that $|0\rangle = \gamma_1|\psi_x\rangle + \gamma_2|\theta_x\rangle$, and $|\gamma_1|^2 = q_x$, and $\langle\theta_x|\psi_x\rangle = 0$. We can apply Fact 4 to A running on state $|\psi_x\rangle$, and get the required number of $x'(i)$, so that the obtained states are close to each other, i.e., E happens with probability $p_x - 2\alpha$. We are interested in the joint probability of events E, F when $U_{x'(i)}$ is performed on $|0\rangle = \gamma_1|\psi_x\rangle + \gamma_2|\theta_x\rangle$. Clearly this probability is at least $q_x(p_x - 2\alpha)$. $\square$

3. A SORTING ALGORITHM

In this section we describe the algorithm needed to prove Theorem 1. For this upper bound we identify time complexity with the number of constant fan-in gates needed to build the circuit (query gates still count as one gate).

The algorithm iterates minimum finding and uses the following result described by Dürr and Høyer [12] based on Grover's famous search algorithm [13].

FACT 6. *There is a quantum query algorithm that, given a comparison oracle to n numbers, finds the minimum of these numbers with probability $1 - \epsilon$, and uses $O(\sqrt{n} \log(1/\epsilon))$ queries (and gates) and space $O(\log^2 n \log(1/\epsilon))$.*

Let S be the space bound. Then $x = x_1, \dots, x_n$ can be partitioned into $b = S/(c \log n)$ blocks y^i of $O(n \log n/S)$ numbers each (for some large enough constant c). Here is the algorithm:

1. FOR $i := 1$ to b compute the position of the minimum of y^i and store it together with i .
2. Arrange the positions of the minima in the memory as a Heap ordered by the minima's size (see [11]).
3. For $i := 1$ to n DO
 - (a) Output (i, j) if the minimal number among the block minima is x_j .
 - (b) Remove j and its block number k from the Heap.
 - (c) Find the position of the minimal number x_l larger than x_j in y^k .
 - (d) Insert (l, k) into the Heap.

Note that all these operations can be performed with a comparison oracle.

Steps 1. and 3.c) employ the minimum finding algorithm of Fact 6. To ensure correctness in step 1. that algorithm is used with error bound $1/S^2$, hence the time for step 1. is $O(S/\log(n) \cdot \sqrt{n \cdot \log(n)/S} \cdot \log S) = O(\sqrt{n \log(n) S})$.

Step 2. can be done in time $O(S)$, steps 3.a) and 3.b) need time $O(\log n)$, and step 3.d) needs time $O(\log S \log n)$ in each iteration [11].

In step 3.c) the algorithm is used with error $1/n^2$. Then the running time for 3.c) is overall $O(n \cdot \sqrt{n \cdot \log(n)/S} \cdot \log n) = O(n^{3/2} \log^{3/2}(n)/\sqrt{S})$.

The time spent in the other steps is dominated by the time used in step 3.c), if $S = O(n/\log n)$.

Note that for reusing the $O(\log^3 n)$ qubits of storage used by the minimum finding algorithm, it is understood that this algorithm is used in the following way. It is run in the usual way, with measurements deferred to the end. Then (instead of measuring) its output is copied to some qubits using controlled-nots. Afterwards the minimum finding algorithm is run *backwards* to clean up the storage it has used. Since the error of the algorithm is small enough this leads to an algorithm with overall error bounded by $O(1/n)$, compare [5, 1] for details on how to run subroutines on a quantum computer.

Also note that the storage bound is not violated. Hence the algorithm behaves as announced.

4. THE LOWER BOUND

Now we give the proof of Theorem 2. After some simple preparations we show how to decompose a quantum circuit into slices that contain only few oracle gates but must (on average) produce many outputs. Then, in our main lemma, we give an upper bound on the number of outputs such a slice can give. In the rest of this section we prove that lemma.

4.1 Preparations

Let A be a quantum circuit with T oracle gates and S work qubits. A contains n output operations, each of which writes on $3 \log n$ of the output qubits. If one of these outputs is measured in the standard basis, then with probability $2/3$

a pair $(\text{Min}(x, j_i), j_i)$ is produced, and with certainty the j_i form a fixed permutation of $\{1, \dots, n\}$.

First we note that the success probability can be improved in some sense.

LEMMA 7. *The success probability can be improved to $1 - \epsilon$ for any constant $\epsilon > 0$ without changing S, T by more than a constant factor, at the expense of adding a circuit consisting of $O(\log n)$ qubits and a majority computation to any output gate.*

PROOF. We can use the circuit some l times “in parallel”. For each output first all l “parallel” outputs are mapped to some extra work storage ($O(l \log n)$ qubits), then an operator is applied to these that computes the most frequent output and maps it to the real output qubits. By standard arguments the error probability drops exponentially in l , and so $l = O(1)$ suffices, hence the increase in space is a factor of l and an additive $O(\log n)$, the increase in time is a factor of l . $\square$

Note that we cannot reuse the $O(\log n)$ qubits, since even if we try to uncompute the computation on them due to the constant error the result will not be close to a blank state. The lower bound proof will enforce the space restriction only between slices of the circuit, so the above construction is still useful.

Due to our output convention now each individual output is correct with probability $1 - \epsilon$ for some arbitrarily small constant ϵ . But then many outputs must be correct simultaneously with high probability.

LEMMA 8. *Assume each individual output is correct with probability $1 - \epsilon$. Let R be any set of l outputs. Then with probability $1 - \sqrt{\epsilon}$ at least $(1 - \sqrt{\epsilon})l$ of the outputs in R are correct.*

PROOF. If the probability that at least $(1 - \sqrt{\epsilon})l$ of the outputs in R are correct simultaneously is less than $1 - \sqrt{\epsilon}$, then the expected number of correct outputs in R is less than $(1 - \sqrt{\epsilon}) \cdot l + \sqrt{\epsilon} \cdot (1 - \sqrt{\epsilon})l = (1 - \epsilon)l$, which is impossible due to the linearity of expectation. $\square$

Hence with large probability at least a big fraction of the outputs are correct.

4.2 Slicing Quantum Circuits

Consider the following way to slice a given quantum circuit A for sorting. Fix some parameter δ . Slice A_1 contains all the gates from the beginning of the circuit up to the $\delta\sqrt{n}$ -th query gate. Slice A_2 contains the next gates until the $2\delta\sqrt{n}$ -th query gate and so on. Overall there are $M = \lceil T/(\delta\sqrt{n}) \rceil$ slices. Note that each slice A_i is a quantum circuit that contains $\delta\sqrt{n}$ queries, and that uses S qubits of work space which are initialized to some state depending on the input being computed by the slices $A_1, A_2, \dots, A_{i-1}$. Note that the average number of outputs in a slice is n/M .

In the following we consider the computational power of an individual slice. We will give an upper bound on the number of outputs a slice can make within the success guarantees derived in the previous subsection. The set of outputs will be restricted to those which output one of the numbers smaller than the median, i.e., outputs for the largest $n/2$ numbers will not be considered. We can now state our main lemma.

LEMMA 9 (MAIN). *Let A be any quantum algorithm that is initially given some S qubits in an arbitrary state depending on a classical input $x = x_1 \dots x_n \in [1, \dots, n^2]^n$, and that afterwards accesses x via $\delta\sqrt{n}$ oracle queries only. Assume that A produces $l \leq n/10$ outputs $O_j(x), \dots, O_{j+l}(x)$, and that for all x and $i \in [j, \dots, j+l]$ the output $O_i(x)$ is correct with probability $1 - \epsilon$. Then for $\delta = 10^{-8}$ it holds that*

$$\begin{aligned} & (1 - \sqrt{\epsilon}) \cdot 2^{-S} \cdot 2^{-l \cdot H(\sqrt{\epsilon})} \\ & \leq (0.99)^{(1 - \sqrt{\epsilon})l/2}. \end{aligned}$$

Note that the space bound enters the main lemma only via the amount of initial information. The function H is the binary entropy function $H(p) = -p \log p - (1 - p) \log(1 - p)$. Let us now deduce Theorem 2 from the lemma.

PROOF OF THEOREM 2: Given is a circuit A . First apply Lemma 7 to reduce the error probability to ϵ for some small enough constant ϵ . Consider any slice A_i of the circuit A and assume that the slice makes some l outputs. If $l = cS$ for some large enough constant c , then the lemma says

$$\begin{aligned} & 1 - \sqrt{\epsilon} \\ & \leq 2^S \cdot 2^{cS \cdot H(\sqrt{\epsilon})} \cdot (0.99)^{(1 - \sqrt{\epsilon})cS/2} < 1/2. \end{aligned}$$

Contradiction! So the number of outputs in the slice is at most $O(S)$.

There are $M \leq \lceil T/(\delta\sqrt{n}) \rceil$ slices producing $n/2$ outputs, but the overall number of outputs is at most $\lceil T/(\delta\sqrt{n}) \rceil \cdot cS$, so $TS = \Omega(n^{3/2})$.

4.3 Proof of the Main Lemma

The plan of the proof is to first show that for each x the expectation over all subsets of $l' = (1 - \sqrt{\epsilon})l$ outputs of the probability that these are simultaneously correct is at least $(1 - \sqrt{\epsilon}) \cdot 2^{-S} \cdot 2^{-l \cdot H(\sqrt{\epsilon})}$, even if the initial information is replaced by a random state.

Afterwards we show that for any l' output positions the expectation over inputs x of the probability that the outputs are simultaneously correct is at most $(0.99)^{l'/2}$ then.

Let $l \leq n/10$. Also fix $\delta = 10^{-8}$.

LEMMA 10. *Suppose an algorithm produces l outputs, and for all inputs x each output $O_i(x)$ is equal to $(\text{Min}(x, j_i), j_i)$ with probability $1 - \epsilon$. Then for all inputs the expectation over sets of l' outputs of the correctness probability is at least $(1 - \sqrt{\epsilon}) \cdot 2^{-H(\sqrt{\epsilon})l}$.*

PROOF. First apply Lemma 8. This lemma implies that with probability $1 - \sqrt{\epsilon}$ at least l' outputs are simultaneously correct.

In other words for all x :

$$\text{Prob}(\exists l' \text{ outputs } O_i(x) = (\text{Min}(x, j_i), j_i)) \geq 1 - \sqrt{\epsilon},$$

where the probability is over the measurements. This implies

$$\sum_{o_1, \dots, o_{l'}} \text{Prob}(\forall i : O_{o_i}(x) = (\text{Min}(x, j_{o_i}), j_{o_i})) \geq 1 - \sqrt{\epsilon}$$

and hence with $\binom{l}{(1 - \sqrt{\epsilon})l} \leq 2^{lH(\sqrt{\epsilon})}$ the result follows. $\square$

Assume that these output gates are supposed to give the outputs $\text{Min}(x, k_1), \dots, \text{Min}(x, k_{l'})$ for some $k_1 < \dots <$

$k_{l'} \leq n/2$ and that they are given by gates $O_1, \dots, O_{l'}$ (renumbered for convenience).

Next we get rid of the initial information, at the expense of increasing the failure probability. We use Lemma 3, restated here in a more specific form.

LEMMA 11. *Suppose there is an algorithm that uses S qubits of initial information and else makes only queries to the input x . Suppose the algorithm outputs a fixed set of input positions $\text{Min}(x, k_1), \dots, \text{Min}(x, k_{l'})$ simultaneously correct with probability P_x .*

Then there is an algorithm that uses no initial information, makes the same number of queries, and has success probability $P_x/2^5$.

At this point we are left with the following problem. We have a circuit that is supposed to output l' numbers from the sorted sequence with some probability P . The circuit accesses the input only via $\delta\sqrt{n}$ queries. We have to show that P is exponentially small in l' . This is established by the following lemma.

LEMMA 12. *For any algorithm that uses $\delta\sqrt{n}$ queries to inputs $x \in [1, \dots, n^2]^n$ and tries to output l' fixed positions of the sorted sequence, $\text{Min}(x, k_1), \dots, \text{Min}(x, k_{l'})$, the expectation over all x of the success probability is at most $0.99^{l'/2}$.*

Note that this lemma together with the previous two lemmas immediately implies the main lemma. In the rest of this section we provide its proof.

Fix some input x . Denote by q_x^i the probability that the outputs $O_1(x)$ up to $O_i(x)$ are correct. Also denote by p_x^i the probability that $O_i(x)$ is correct conditioned on the event that the previous outputs are correct. Then $q_x^{i+1} = q_x^i p_x^{i+1}$. We are interested in bounding $E[q_x^{l'}]$. We plan to show that for most i there is a constant factor gap between $E[q_x^i]$ and $E[q_x^{i+1}]$.

The next proposition states the main adversary argument.

PROPOSITION 1. *Let x and i be given. Then there is a set $J_{x,i} \subseteq \{1, \dots, n\}$ of size $n/2 - \delta n$ so that for each $j \in J_{x,i}$ there are*

$$|\{\text{Min}(x, k_{i+1} - 1), \dots, \text{Min}(x, k_{i+1}) - 1\}|$$

inputs $x'(j)$ so that with probability $q_x^i(p_x^{i+1} - 2\sqrt{\delta})$ the first i outputs are correct and the $i+1$ st is incorrect on $x'(j)$.

Furthermore, each such $x'(j)$ is "generated" in this way by at most n^2 inputs x .

PROOF. We want to apply Lemma 5. Let the condition F of that lemma be that the outcome of the first i measurements is correct. The event E is set to be that the $i+1$ st measurement is correct, i.e., equals $\text{Min}(x, k_{i+1})$. Clearly F occurs with probability q_x^i on x and E occurs conditionally with probability p_x^{i+1} . Then the lemma tells us that we may switch $(1 - \delta)n$ positions in x arbitrarily, and still get the same measurement results with probability $q_x^i(p_x^{i+1} - 2\sqrt{\delta})$. To avoid changing the correctness of previous outputs we only flip those positions containing numbers larger than $\text{Min}(x, k_{i+1})$ (which is at most the median). Thus we can flip at least $n/2 - \delta n$ positions.

If we switch a position in x so that $x'(j)$ has a different $\text{Min}(x, k_{i+1})$, then we create an error in the $i+1$ th output on $x'(j)$. Note that if we change an input position j in x , so that

$\text{Min}(x, k_i) \leq \text{Min}(x, k_{i+1} - 1) \leq x'(j)_j < \text{Min}(x, k_{i+1})$, then if the k_{i+1} -th output on x is correct, the k_{i+1} -th output on $x'(j)$ is wrong, and all previous outputs stay correct.

Also note that each $x'(j)$ is derived from at most n^2 inputs x , since to change $x'(j)$ to x we have to change $\text{Min}(x'(j), k_{i+1})$ to some other value. $\square$

Now it is clearly true that

$$E[q_x^i] = E[q_x^i \cdot (p_x^{i+1} + 1 - p_x^{i+1})],$$

and $q_x^i(1 - p_x^{i+1})$ is the probability that the first i outputs are correct and the $i+1$ st is wrong.

$$\begin{aligned} & q_x^i(1 - p_x^{i+1}) \\ & \geq \max_{y: x=y'(j)} q_y^i(p_y^{i+1} - 2\sqrt{\delta}) \\ & \geq \sum_{y: x=y'(j)} (q_y^i(p_y^{i+1} - 2\sqrt{\delta}))/n^2 \end{aligned}$$

due to the above proposition, where the notation $x = y'(j)$ indicates that x can be derived by a change of one position inserting small error as described above. Since we can redistribute the terms in the expectation we get

$$\begin{aligned} & E[q_x^i] \\ & \geq E[q_x^i p_x^{i+1} \\ & + q_x^i(p_x^{i+1} - 2\sqrt{\delta}) \cdot (n/2 - \delta n) \\ & \cdot \frac{|\{\text{Min}(x, k_{i+1} - 1), \dots, \text{Min}(x, k_{i+1})\}| - 1}{n^2}]. \end{aligned}$$

Therefore, denoting

$$D(x, i+1) = (|\{\text{Min}(x, k_{i+1} - 1), \dots, \text{Min}(x, k_{i+1})\}| - 1)/n :$$

$$\begin{aligned} & E[q_x^i] + \sqrt{\delta} E[q_x^i \cdot D(x, i+1)] \\ & \geq E[q_x^{i+1}] + (1/2 - \delta) \cdot E[q_x^{i+1} \cdot D(x, i+1)]. \end{aligned}$$

Hence there is a constant factor gap between $E[q_x^i]$ and $E[q_x^{i+1}]$, unless $E[q_x^i \cdot D(x, i+1)]$ is much larger than $E[q_x^i]$ or $E[q_x^{i+1} \cdot D(x, i+1)]$ is much smaller than $E[q_x^{i+1}]$. For each position where this is not true we can collect a constant factor.

Intuitively the success probability may well change with smaller or larger $D(x, i)$. But if this happens too often, the algorithm will concentrate too much on inputs with many small (or many large) $D(x, i)$, also decreasing the success probability. We will show that in this case the success probability is exponentially small in l' . We will show this only for the second case of small $D(x, i)$, the other case is similarly.

We call an output position i *bad* (for short), if

$$E[q_x^i D(x, i)] < E[q_x^i]/32,$$

and *bad* (for long), if

$$E[q_x^{i-1} D(x, i)] > 32E[q_x^{i-1}].$$

Clearly for each i that is not bad (neither for short nor long)

$$E[q_x^i] \leq E[q_x^{i-1}] \cdot .99.$$

Hence we are interested in the situation where there are $l'/4$ bad i (for short), since otherwise (with less than $l'/4$ bad i for long and short intervals each) Lemma 12 is proved.

Note that for each bad i the contribution of those x with $D(x, i) < 1/4$ to $E[q_x^i]$ is at least $(1 - 1/8)E[q_x^i]$.

Consider a matrix with rows labeled by inputs x and columns by elements of $\{1, \dots, l'/4\}$. We investigate how well we can put weights q_x^i on this matrix according to the following rules (call all x, i with $D(x, i) \leq 1/4$ *short*):

1. $q_x^i \leq 1$ for all x, i .
2. $q_x^i \leq q_x^{i-1}$ for all short x, i .
3. $\sum_{x: D(x, i) < 1/4} \frac{1}{n^{2\pi}} q_x^i \geq 7/8 \cdot E[q_x^i]$.

We are interested in optimizing $E[q_x^{l'/4}]$. Note that we may neglect the i which are not bad. We view this as a game played on the columns of the matrix between a player who is allowed to distribute weight on the x, i , and a player, who can choose the order of the columns.

PROPOSITION 2. *If there are $l'/4$ bad i , then $E[q_x^{l'/4}] \leq (1/2)^{l'/8}$. In particular Lemma 12 holds.*

PROOF. As said above, this is shown by playing a game between Player 1, who distributes weight on the next column, and Player 2, who picks the next column (i.e., the next i), which will be done randomly. We will show that with high probability the expected weight drops by $1/2$ at least in each round.

It is not hard to see that

$$\text{Prob}(D(x, i) < 1/4 \mid D(x, i_1), \dots, D(x, i_t) < 1/4) \leq 1/4,$$

for all such conditions, since the probability that one interval is short does rather decrease the probability that another interval is short.

The starting column i_1 is arbitrary. In the beginning Player 1 obviously sets all $q_x^{i_1} = 1$ for the short x, i_1 , and then distributes weight $1/4 \cdot 1/7$ to the remaining x, i_1 . In each following round Player 1 will also set each $q_x^{i_j}$ for short x, i_j equal to its predecessor, and distribute $1/7$ th of the obtained overall weight in the column to the other x, i_j .

Assume that Player 1 always distributes weight uniformly, where she has a choice. Then $E[q_x^{i_1}] \leq 1/4 \cdot 8/7 = 2/7$, and in general

$$E[q_x^{i_j}] \leq \sum_{k=0}^{j-1} E[q_x^{i_k}] / 7 \cdot 4^{k-j},$$

setting $E[q_x^{i_0}] = 8$ for convenience. It is not hard to show that $E[q_x^{i_j}] < 1/2^j$ in this situation.

Now consider Player 1's options to distribute weights non-uniformly. Since in the first $l'/8$ rounds there are always at least $l'/8$ positions left, it can be shown that if $E[q_x^{i_j}]$ is not already small, for most positions Player 2 picks, Player 1's choice is not much better than uniform. To do so one uses that the events " x, i short" (resp. the corresponding sets of inputs) behave like independent events, and if $E[q_x^{i_j}]$ is bigger than $2^{-l'/\text{const}}$, then for most of them not much more weight than under the uniform distribution can be gathered. We omit the details of this argument. $\square$

5. CONCLUSIONS AND OPEN PROBLEMS

We have shown that quantum computers, though in general not much faster for the important task of sorting, outperform classical computers significantly in space bounded

sorting. This setting is motivated by applications with distributed datasets too large to fit into the working memory of a single computer. Furthermore understanding the complexity of such a basic problem is important in itself. We have seen that for sorting a tradeoff result of the form $T^2S \leq \tilde{O}(n^3)$ exists, as opposed to the classical tradeoff $TS = \Theta(n^2)$. Exploring the question whether the upper bound on quantum sorting is tight we have proved a lower bound of $TS = \Omega(n^{3/2})$, showing that for small space bounds the algorithm is not too far from optimal.

Our result can easily be extended to an $ST = \Omega(n^2)$ lower bound for classical sorting in the situation that all input positions are known to hold mutually distinct numbers, by considering circuit slices of length δn .

The most important open problem is, whether the lower bound can be improved to (almost) match the upper bound, which is what we conjecture. To prove such a result it seems one should show that circuit slices containing $\sqrt{nl}$ queries cannot produce more than $O(l)$ outputs. To do so it would suffice to show that any quantum algorithm with $\delta\sqrt{nl}$ queries that tries to compute l elements of the sorted sequence succeeds only with probability $2^{-\Theta(l)}$.

It is also of interest what the time-space complexity of sorting is if we allow no errors. In this situation approaches based on Grover search fail and it is well possible that the same tradeoff as classically holds.

Consider the element distinctness problem, i.e., deciding whether n given numbers are all pairwise different. It is conjectured that for classical computers the element distinctness problem has about the same complexity as sorting, and is thus a decision problem capturing the difficulty of sorting. Buhrman et al. describe in [9] a quantum algorithm that runs in sublinear time, and a slight variation of their algorithm achieves a tradeoff of $T^2S = \tilde{O}(n^2)$ for deciding element distinctness. Hence in the quantum case element distinctness is strictly easier than sorting due to the lower bound for sorting given in this paper, consider e.g. the case that $S \leq \text{poly}(\log n)$. Can a matching lower bound be shown for element distinctness? Is element distinctness really as hard as sorting classically? Note that strong classical tradeoffs are known for element distinctness if only a comparison oracle is used [8, 19], but the best tradeoff known for general models is $T = \Omega(n \log(n/S))$, given in [4]. A quantum query lower bound of $\Omega(n^{2/3})$ has recently been shown by Shi [18].

Another open problem concerns the query problem we have analyzed. We have shown that finding l elements of the sorted sequence within $\delta\sqrt{n}$ queries is possible with exponentially small success probability only. More generally, suppose one is given l instances of a query problem, and is allowed to do a number q of queries that is known to give the result for one instance with probability $p < 1$ only, on average over all inputs. How large is the success probability of computing correctly on l instances? In other words, does a direct product result hold for quantum black-box algorithms?

Acknowledgements

The author wishes to thank Peter Høyer, Alexander Razborov, Avi Wigderson, Ronald de Wolf, and the anonymous referees for helpful comments and discussions.

6. REFERENCES

- [1] D. Aharonov, A. Kitaev, and N. Nisan. Quantum circuits with mixed states. *30th ACM Symposium on Theory of Computing*, pp. 20–30, 1998. Also: [quant-ph/9806029](#).
- [2] A. Ambainis. Quantum lower bounds by quantum arguments. *32nd ACM Symposium on Theory of Computing*, pp. 636–643, 2000. Also: [quant-ph/0002066](#).
- [3] P. Beame. A general sequential time-space tradeoff for finding unique elements. *SIAM Journal on Computing*, vol.20, pp.270–277, 1991.
- [4] P. Beame, M. Saks, X. Sun, E. Vee. Super-linear time-space tradeoff lower bounds for randomized computation. *41st IEEE Symposium on Foundations of Computer Science*, pp.169–179, 2000.
- [5] C.H. Bennett, E. Bernstein, G. Brassard, and U. Vazirani. Strengths and weaknesses of quantum computing. *SIAM Journal on Computing*, vol. 26, pp. 1510–1523, 1997. Also: [quant-ph/9701001](#).
- [6] C.H. Bennett, G. Brassard, C. Crepeau, R. Josza, A. Peres, W. Wootters. Teleporting an Unknown Quantum State via Dual Classical and Einstein-Podolsky-Rosen Channels. *Phys. Rev. Lett.*, vol.70, pp.1895–1899, 1993.
- [7] A. Borodin, S. Cook. A time-space tradeoff for sorting on a general sequential model of computation. *SIAM Journal on Computing*, vol.11, pp.287–297, 1982.
- [8] A. Borodin, F. Fich, F. Meyer and der Heide, E. Upfal, A. Wigderson. A time-space tradeoff for element distinctness. *SIAM Journal on Computing*, vol.16, pp.97–99, 1987.
- [9] H. Buhrman, C. Dürr, M. Heiligman, P. Høyer, F. Magniez, M. Santha, R. de Wolf. Quantum Algorithms for Element Distinctness. *IEEE Conference on Computational Complexity*, pp.120–130, 2001.
- [10] H. Buhrman, R. de Wolf. Complexity Measures and Decision Tree Complexity: A Survey. To appear in *Theoretical Computer Science*, 2002.
- [11] T.H. Cormen, C.E. Leiserson, R.L. Rivest, C. Stein. Introduction to Algorithms. MIT Press, 2001.
- [12] C. Dürr, P. Høyer. A quantum algorithm for finding the minimum. [quant-ph/9607014](#), 1996.
- [13] L.K. Grover. A fast quantum mechanical algorithm for database search. *28th ACM Symposium on Theory of Computing*, pp. 212–219, 1996. Also: [quant-ph/9605043](#).
- [14] P. Høyer, J. Neerbek, Y. Shi. Quantum complexities of ordered searching, sorting, and element distinctness. *28th International Colloquium on Automata, Languages, and Programming*, pp.62–73, 2001. Also: [quant-ph/0102078](#).
- [15] M.A. Nielsen and I.L. Chuang. Quantum Computation and Quantum Information. Cambridge University Press, 2000.
- [16] J. Pagter, T. Rauhe. Optimal Time-Space Trade-Offs for Sorting. *39th IEEE Symposium on Foundations of Computer Science*, pp.264–268, 1998.
- [17] S. Reisch, G. Schnitger. Three Applications of Kolmogorov-Complexity. *23rd IEEE Symposium on Foundations of Computer Science*, pp.45–52, 1982.
- [18] Y. Shi. Quantum Lower Bounds for the Collision and the Element Distinctness Problems. *43rd IEEE Symposium on Foundations of Computer Science*, pp.513–519, 2002.
- [19] A.C.C. Yao. Near optimal time-space tradeoffs for element distinctness. *29th IEEE Symposium on Foundations of Computer Science*, pp.91–97, 1988.